ISCR

Institute for Scientific Computing Research

Scalable Linear Solvers Workshop Report

CASC

Center for Applied Scientific Computing

Working to Solve Large-Scale Linear Systems

Robert D. Falgout and Jinchao Xu

Synopsis of Workshop Events

The multigrid method has been actively investigated by researchers for many years, but how can it be made more practical in applications? This question was discussed in detail at a workshop sponsored by CASC, the Center for Applied Scientific Computing at Lawrence Livermore National Laboratory.

This three-day (June 23–25, 1999) Oberwolfach-style workshop was held in the beautiful Wente Vineyards of Livermore, California. Just a few minutes' drive from CASC's office, the informal setting was ideal for the size and format of this workshop. About 40 invited participants attended (mostly from the United States, with a few from Europe) and joined in lively discussions during and after the workshop sessions.

This workshop was one aspect of the many research outreach activities sponsored by CASC and the Institute for Scientific Computing Research (ISCR). Under the direction of Steven Ashby, CASC has grown from 12 scientists in 1996 to more than 80 at present. CASC conducts collaborative scientific investigations that require the power of high-performance computers and the efficiency of modern computational methods. Its research and development activities are applications-driven, and focus on those LLNL programmatic objectives requiring advanced computational technologies. The Center's core competencies include high-performance computing, computational physics, numerical mathematics, and computer science.

The workshop was organized by researchers in the Scalable Linear Solvers project in CASC (http://www.llnl.gov/CASC/linear solvers/). The project develops scalable algorithms and software for solving large, sparse linear systems of equations on parallel computers. A major research focus is multigrid,one of the most powerful solution methodologies for solving large-scale systems. Ashby and his colleagues in CASC have placed the multigrid method at the center of their algorithmic research development. They intend to accelerate and broaden its application, which will have a notable impact on scientific simulation.

The multigrid method has been shown to be extremely efficient both in theory and in practice. Traditionally, however, it has been somewhat problem-dependent and difficult to use; particularly in a software library setting. In recent years, there has been a desire to make the method more accessible to general users. This has led to a resurgence of interest in a class of algorithms originally developed in the early Eighties, algebraic multigrid methods.

Similar to classical multigrid methods, algebraic multigrid methods exhibit optimal computational complexities for many applications, but are relatively much easier to use (as are other algebraic algorithms such as Gaussian elimination). These algorithms were the major focus of the workshop. In CASC, a large group of researchers works on this subject, including Andy Cleary, Robert Falgout, Van Henson, Jim Jones, Barry Lee, Beth Ong, Charles Tong, Panayot Vassilevski, and Ulrike Yang. The group has formed several collaborations, including a close partnership with University of Colorado at Boulder researchers Tom Manteuffel, Steve McCormick, and John Ruge. As the participants are mostly very active researchers on the subject, the atmosphere of the workshop was very lively and down-to-earth.

The workshop began at 9 a.m. on June 23, 1999, with Ashby's opening remarks. Falgout then explained the meeting format and program, one that encouraged an atmosphere of open exchange. Each morning and after-

Working to Solve Large-Scale Linear Systems (continued)

noon session was devoted to a special topic. Sessions opened with a few short presentations to enumerate the issues, then lively discussions followed.

The first session of the workshop, organized by Raytcho Lazarov, related to domain decomposition techniques, and focused primarily on methods for non-matching grids. Three sessions were devoted to algebraic multigrid methods. Organized by Jim Jones and Panayot Vassilevski from LLNL, two of the sessions centered around recent algorithmic advances for improving robustness. Van Henson from LLNL organized the third session, which covered the important and difficult problem of developing parallel algebraic multigrid algorithms. Another session, organized by Falgout, discussed issues related to implementing and designing multigrid methods for today's high-performance computers, where parallelism and cache-performance play a dominant role. Finally, Edmond Chow from LLNL organized a session on multilevel ILU, a method that has much in common with algebraic multigrid, but one that brings a different algorithmic perspective to the table.

In addition to the special sessions, another major activity was the afternoon group discussions. During the lunch breaks, participants could either stick together to form their own small work groups or join an organized discussion. As it turned out, most participants joined the organized sessions, which were meant to be informal and deliberately provocative.

The first such discussion was organized by Jinchao Xu, and asked the question, “Which problems can and cannot be solved with multigrid, and why is multigrid not commonly used in practice?” As expected, this generated much controversy. In fact, the question itself was the first thing to be debated, as some of the participants agreed with the statement that multigrid was not commonly used in practice, and others disagreed. Xu constructed a web page (using his laptop) with a list of important problems, and solicited input on those problems for which multigrid is an effective solver. Again, much debate ensued, and although the details could not quite be agreed upon, an overall consensus emerged that multigrid methods can be effective solvers for a remarkably wide range of problems. This discussion carried over to the next lunch.

The second afternoon discussion occurred on the third day of the meeting and was organized by Tom Manteuffel. The similarities and differences between multigrid, domain decomposition, and aggregation methods were explored, and the question was asked, “Can one establish a common framework into which these all fit?” Manteuffel outlined a framework in which methods are viewed as “approximate” Schur complement methods. This also generated much healthy discussion, and a relevant early framework based on space decomposition and subspace correction was mentioned.

On the basis of the productivity of the workshop and the enthusiasm of its participants, CASC plans to sponsor three or four workshops per year on subjects related to its core mission. For fiscal year 2000, these will include scalable nonlinear solvers, transport methods, and mining large scientific data sets. For information on future workshops, watch the webpage at http://www.llnl.gov/casc/workshops.

University of California
Lawrence Livermore National Laboratory
Technical Information Department
Livermore, CA 94551